

Internet Engineering Task Force (IETF)
Request for Comments: 9081
Updates: 6514
Category: Standards Track
ISSN: 2070-1721

Z. Zhang
L. Giuliano
Juniper Networks
July 2021

Interoperation between Multicast Virtual Private Network (MVPN) and Multicast Source Directory Protocol (MSDP) Source-Active Routes

Abstract

This document specifies the procedures for interoperation between Multicast Virtual Private Network (MVPN) Source-Active (SA) routes and customer Multicast Source Discovery Protocol (MSDP) SA routes, which is useful for MVPN provider networks offering services to customers with an existing MSDP infrastructure. Without the procedures described in this document, VPN-specific MSDP sessions are required among the Provider Edge (PE) routers that are customer MSDP peers. This document updates RFC 6514.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9081>.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
 - 1.1. MVPN RPT-SPT Mode
2. Terminology
 - 2.1. Requirements Language
3. Specification
4. Security Considerations
5. IANA Considerations
6. References
 - 6.1. Normative References
 - 6.2. Informative References

Acknowledgements

Authors' Addresses

1. Introduction

Section 14 ("Supporting PIM-SM without Inter-Site Shared C-Trees") of [RFC6514] specifies the procedures for MVPN PEs to discover (C-S,C-G) via MVPN Source-Active A-D routes and then send Source Tree Join (C-S,C-G) C-multicast routes towards the ingress PEs to establish shortest path trees (SPTs) for customer Any-Source Multicast (ASM) flows for which they have downstream receivers. (C-*,C-G) C-multicast routes are not sent among the PEs, so inter-site shared C-Trees are not used, and the method is generally referred to as "spt-only" mode.

With this mode, the MVPN Source-Active routes are functionally similar to MSDP Source-Active messages. For a VPN, one or more of the PEs, say PE1, either acts as a C-RP and learns of (C-S,C-G) via PIM Register messages or has MSDP sessions with some MSDP peers and learns of (C-S,C-G) via MSDP SA messages. In either case, PE1 will then originate MVPN SA routes for other PEs to learn (C-S,C-G).

[RFC6514] only specifies that a PE receiving the MVPN SA routes, say PE2, will advertise Source Tree Join (C-S,C-G) C-multicast routes if it has corresponding (C-*,C-G) state learnt from its Customer Edge (CE). PE2 may also have MSDP sessions for the VPN with other C-RPs at its site, but [RFC6514] does not specify that PE2 advertises MSDP SA messages to those MSDP peers for the (C-S,C-G) that it learns via MVPN SA routes. PE2 would need to have an MSDP session with PE1 (that advertised the MVPN SA messages) to learn the sources via MSDP SA messages for it to advertise the MSDP SA to its local peers. To make things worse, unless blocked by policy control, PE2 would in turn advertise MVPN SA routes because of those MSDP SA messages that it receives from PE1, which are redundant and unnecessary. Also notice that the PE1-PE2 MSDP session is VPN specific (i.e., only for a single VPN), while the BGP sessions over which the MVPN routes are advertised are not.

If a PE does advertise MSDP SA messages based on received MVPN SA routes, the VPN-specific MSDP sessions with other PEs are no longer needed. Additionally, this MVPN/MSDP SA interoperation has the following inherent benefits for a BGP-based solution.

- * MSDP SA refreshes are replaced with BGP hard state.
- * Route reflectors can be used instead of having peer-to-peer sessions.
- * VPN extranet [RFC2764] mechanisms can be used to propagate (C-S,C-G) information across VPNs with flexible policy control.

While MSDP Source-Active routes contain the source, group, and RP addresses of a given multicast flow, MVPN Source-Active routes only contain the source and group. MSDP requires the RP address information in order to perform MSDP peer Reverse Path Forwarding (RPF). Therefore, this document describes how to convey the RP address information into the MVPN Source-Active route using an Extended Community so this information can be shared with an existing MSDP infrastructure.

The procedures apply to Global Table Multicast (GTM) [RFC7716] as well.

1.1. MVPN RPT-SPT Mode

For comparison, another method of supporting customer ASM is generally referred to as "rpt-spt" mode. Section 13 ("Switching from a Shared C-Tree to a Source C-Tree") of [RFC6514] specifies the MVPN SA procedures for that mode, but those SA routes are a replacement for PIM-ASM assert and (s,g,rpt) prune mechanisms, not for source discovery purposes. MVPN/MSDP SA interoperation for the "rpt-spt" mode is outside the scope of this document. In the rest of the document, the "spt-only" mode is assumed.

2. Terminology

Familiarity with MVPN [RFC6513] [RFC6514] and MSDP [RFC3618] protocols and procedures is assumed. Some terminology is listed below for convenience.

ASM: Any-Source Multicast

SPT: source-specific Shortest Path Tree

RPT: Rendezvous Point Tree

C-S: a multicast source address, identifying a multicast source located at a VPN customer site

C-G: a multicast group address used by a VPN customer

C-RP: a multicast Rendezvous Point for a VPN customer

C-multicast: a multicast for a VPN customer

EC: Extended Community

GTM: Global Table Multicast, i.e., a multicast in the default or global routing table vs. a VPN Routing and Forwarding (VRF) table

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Specification

The MVPN PEs that act as customer RPs or have one or more MSDP sessions in a VPN (or the global table in case of GTM) are treated as an MSDP mesh group for that VPN (or the global table). In the rest of the document, it is referred to as the PE mesh group. This PE mesh group MUST NOT include other MSDP speakers and is integrated into the rest of the MSDP infrastructure for the VPN (or the global table) following normal MSDP rules and practices.

When an MVPN PE advertises an MVPN SA route following procedures in [RFC6514] for the "spt-only" mode, it MUST attach an "MVPN SA RP-address Extended Community". This is a Transitive IPv4-Address-Specific Extended Community. The Local Administrator field is set to zero, and the Global Administrator field is set to an RP address determined as the following:

- * If the (C-S,C-G) is learnt as a result of the PIM Register mechanism, the local RP address for the C-G is used.
- * If the (C-S,C-G) is learnt as a result of incoming MSDP SA messages, the RP address in the selected MSDP SA message is used.

In addition to the procedures in [RFC6514], an MVPN PE may be provisioned to generate MSDP SA messages from received MVPN SA routes, with or without local policy control. If a received MVPN SA route triggers an MSDP SA message, the MVPN SA route is treated as if a corresponding MSDP SA message was received from within the PE mesh group and normal MSDP procedure is followed (e.g., an MSDP SA message is advertised to other MSDP peers outside the PE mesh group). The (S,G) information comes from the (C-S,C-G) encoding in the MVPN SA Network Layer Reachability Information (NLRI), and the RP address comes from the "MVPN SA RP-address EC" mentioned above. If the received MVPN SA route does not have the EC (this could be from a legacy PE that does not have the capability to attach the EC), the local RP address for the C-G is used. In that case, it is possible that the RP inserted into the MSDP SA message for the C-G is actually the MSDP peer to which the generated MSDP message is advertised,

causing the peer to discard it due to RPF failure. To get around that problem, the peer SHOULD use local policy to accept the MSDP SA message.

An MVPN PE MAY treat only the best MVPN SA route selected by the BGP route selection process (instead of all MVPN SA routes) for a given (C-S,C-G) as a received MSDP SA message (and advertise the corresponding MSDP message). In that case, if the selected best MVPN SA route does not have the "MVPN SA RP-address EC" but another route for the same (C-S, C-G) does, then the next best route with the EC SHOULD be chosen. As a result, if/when the best MVPN SA route with the EC changes, a new MSDP SA message is advertised if the RP address determined according to the newly selected MVPN SA route is different from before. The MSDP SA state associated with the previously advertised MSDP SA message with the older RP address will be timed out.

4. Security Considerations

[RFC6514] specifies the procedure for a PE to generate an MVPN SA upon discovering a (C-S,C-G) flow (e.g., via a received MSDP SA message) in a VPN. This document extends this capability in the reverse direction -- upon receiving an MVPN SA route in a VPN, generate a corresponding MSDP SA and advertise it to MSDP peers in the same VPN. As such, the capabilities specified in this document introduce no additional security considerations beyond those already specified in [RFC6514] and [RFC3618]. Moreover, the capabilities specified in this document actually eliminate the control message amplification that exists today where VPN-specific MSDP sessions are required among the PEs that are customer MSDP peers, which lead to redundant messages (MSDP SAs and MVPN SAs) being carried in parallel between PEs.

5. IANA Considerations

IANA registered the following in the "Transitive IPv4-Address-Specific Extended Community Sub-Types" registry:

Value	Description
0x20	MVPN SA RP-address Extended Community

Table 1

6. References

6.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC3618] Fenner, B., Ed. and D. Meyer, Ed., "Multicast Source Discovery Protocol (MSDP)", RFC 3618, DOI 10.17487/RFC3618, October 2003, <<https://www.rfc-editor.org/info/rfc3618>>.

[RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<https://www.rfc-editor.org/info/rfc6514>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

6.2. Informative References

- [RFC2764] Gleeson, B., Lin, A., Heinanen, J., Armitage, G., and A. Malis, "A Framework for IP Based Virtual Private Networks", RFC 2764, DOI 10.17487/RFC2764, February 2000, <<https://www.rfc-editor.org/info/rfc2764>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.
- [RFC7716] Zhang, J., Giuliano, L., Rosen, E., Ed., Subramanian, K., and D. Pacella, "Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures", RFC 7716, DOI 10.17487/RFC7716, December 2015, <<https://www.rfc-editor.org/info/rfc7716>>.

Acknowledgements

The authors thank Eric Rosen, Vinod Kumar, Yajun Liu, Stig Venaas, Mankamana Mishra, Gyan Mishra, Qin Wu, and Jia He for their reviews, comments, questions, and suggestions for this document.

Authors' Addresses

Zhaohui Zhang
Juniper Networks

Email: zhang@juniper.net

Lenny Giuliano
Juniper Networks

Email: lenny@juniper.net